

5

TITLE OF THE INVENTION  
METHOD AND APPARATUS FOR DISTRIBUTED DATA ARCHIVING

CROSS REFERENCE TO RELATED APPLICATIONS

This application claims priority under 35 U.S.C.  
10 §119(e) to provisional patent application serial number  
60/228,631 filed August 29, 2000, the disclosure of which  
is hereby incorporated by reference.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR  
15 DEVELOPMENT  
N/A

BACKGROUND OF THE INVENTION

20 The diagnosis and treatment of cardiac disease  
requires the acquisition of large amounts of data about  
the patient's medical condition. This data typically  
consists of medical images, patient information and  
clinical reports. The treatments of a single patient's  
disease can generate 1,000,000,000 bytes (1 Gbyte) or  
25 more of digital information. Millions of patients are  
treated worldwide each year for cardiac disease. Many  
other medical conditions exist for which similar  
quantities of data are maintained. The storage and  
management of such patient data requires more capacity  
30 than can be cost effectively deployed using standard  
magnetic disk technology.

1

ATTORNEY DOCKET NO. HEART-001XX  
WEINGARTEN, SCHURGIN,  
GAGNEBIN & HAYES LLP  
TEL. (617) 542-2290  
FAX. (617) 451-0313

Express Mail Number  
EL751780463US

The standard industry solution to this problem has been to combine small amounts of hard disk storage with large amounts of low cost magnetic tape storage. This combination of storage technology is implemented as a storage hierarchy. The premise of such a system is that newer data is accessed more frequently than old data. The typical Hierarchical Storage Management (HSM) system software used to manage a storage hierarchy moves data from the more costly hard disk storage to tape storage based on heuristics such as age of the data or last access time. HSM supports this type of data flow transparently to the software application that generates the data. In other words, if the data were a word processed document called mydoc.doc, access to a document is achieved by opening the file mydoc.doc. The HSM system provides access to data by restoring the data from tape to disk automatically. This transparency is created by a virtual file system (VFS). The VFS concatenates all of the storage in the hierarchy into one logical disk drive. The only apparent difference to the user is speed of access; files on the hard disk are accessed very quickly while files on the tape library are accessed with lengthy delays.

The transparency created by HSM creates problems due to monolithic access, proprietary formats, replication difficulties, and poor performance which make HSM poorly suited to mission critical medical image storage environments. Specifically, the primary design principle of HSM is to create transparency so that existing software does not need to be rewritten to take advantage

of the large storage system. The transparency is achieved by aggregating all individual storage units (single tapes in a robotic library system plus hard disk(s)) into a single file system. As a result, files will be spread across all of the storage units in the hierarchy.

This design creates a situation where an individual storage unit is only meaningful in the context of the HSM system, which created it. Therefore, if a tape is removed from an HSM managed library, the tape will have no meaning outside of this HSM managed library and the tape must be restored to the HSM system to be accessed. HSM manufacturers often use proprietary or limited logical formats on each storage unit so that tapes cannot be read without the HSM software.

Although HSM often has a system for routine backup, HSM is viewed as a separate function from the backup and recovery management needed in the event of a catastrophe. As a result, HSM systems fail to provide for the creation of offsite copies of tapes for efficiently recreating the HSM in the event of a disaster. The complex data flow in HSM that moves files back and forth between levels in the hierarchy results in poor performance. An application must wait for a file to be fully restored from tape to disk before accessing the first byte of information.

## BRIEF SUMMARY OF THE INVENTION

The archive system according to the present invention provides long term storage for large amounts of data. This system is particularly suited for multi-year storage of medical data such as cardiac images, patient demographics and reports. In a preferred embodiment, the archive system of the present disclosure employs one or more digital versatile disks (DVDs) as the storage media. In the archival process, all of the data for a particular patient, procedure or study is stored within one DVD and each DVD is uniquely identified. Also, each DVD may have an executable program stored thereon for independently accessing the archived data from the selected DVD.

Before all of the data is archived, the present archive system segments the data into a plurality of information groups where each group is based on data for the particular patient, procedure or study. Once the information groups are formed, each group is stored on a storage medium (DVD). Before recording an information group on a DVD, the DVD is reviewed to determine whether a sufficient amount of storage space is present to ensure that the information group will be entirely stored within one DVD. Once a DVD is found with enough storage capacity, the information group is stored thereon. This process is repeated for each information group. Although other storage media may be used, DVDs are preferable because their storage capacity is sufficient to store at least one typical information group.

In addition to the one or more information groups stored on each DVD, the unique identification and executable programs stored thereon allow each DVD or subsets of DVDs to be accessed independent from the other DVDs for viewing and printing the data on other processors or work stations. Also, the data from DVDs or DVD subsets created by different processors may be merged and modified to create new information groups.

#### BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWING

The invention will be more fully understood by reference to the following detailed description of the invention in conjunction with the drawings, of which:

Figs. 1(a)-1(d) illustrate configurations for the archival storage media according to embodiments of the present invention;

Fig. 2 illustrates an archival system according to an embodiment of the present invention;

Fig. 3 is a flow chart of an archive session according to an embodiment of the present invention; and

Fig. 4 is a flow chart of an archive selection according to an embodiment of the present invention.

#### DETAILED DESCRIPTION OF THE INVENTION

To provide the storage of large amounts of data in the archive system according to the embodiments of the present invention, the data is segmented into information

groups and stored onto archival storage media. For instance, data may be accessed from one or more external source(s) and segmented into information groups based on a desired procedure, study or patient.

5 Figs. 1(a)-1(b) illustrate configurations for storing such information groups in archival storage media by the present archive system. In Fig. 1(a), retrieved data segmented into information groups 20A, 20B, 20C, 20D, and 20E are recorded onto archival storage media 10 and 12. Four of the information groups 20A, 20B, 20C, and 20D and an index file 14 including an identifier, a database file, a data program, an image player, and/or other programs files are recorded on the archival storage medium 10 before a predetermined storage capacity is reached on the medium 10. The index file 14 allows the information groups 20A, 20B, 20C, and 20D stored on the archival storage medium 10 to be accessed, processed and viewed on a wide range of personal computers and operating systems.

20 The archival storage medium 10 is preferably a DVD, which is capable of storing one or more information groups. However, other archival storage media may be used as long as sufficient capacity is provided for storing at least one information group.

25 Before recording an information group, the recording processor determines whether the entire information group is able to be stored on the medium without exceeding its storage capacity. In the present illustrative example of Fig. 1(a), the four information groups 20A, 20B, 20C, and 30 20D are recorded onto the archival storage medium 10

without exceeding its storage capacity. In this example, the information group 20E cannot be recorded on the archival storage medium 10 without exceeding its storage capacity. Therefore, the information group 20E and an index file 16 including an identifier, a database file, a data program, an image player, and/or other programs files are recorded on the archival storage medium 12 as long as its storage capacity is not exceeded. The index file 16 allows the information group 20E to be accessed, processed, and viewed independently. This system and storage configuration ensures that each information group and the necessary programs and files to retrieve and view the respective information are entirely stored within one archival storage medium. As a result, each archival storage medium stands alone and is capable of being utilized independent from the other archival storage media so that its information can be accessed, viewed and processed.

The present archive system also allows for information and subsets of information stored on one or more disks to be accessed and merged to form new information groups. For instance, the information group 20C from the archival storage medium 10 can be accessed and recorded onto another archival storage medium 30 as illustrated in Fig. 1(b). Similar to the above described archival storage media, an index file 32 is recorded on the archival storage medium 30 to include an identifier, a database file, a data program, an image player, and/or other programs files for accessing, processing and viewing the information group 20C.

Fig. 2 illustrates an archive system according to one embodiment of the present invention. A user interface (UI) 200 controls the present system for retrieving or receiving information from one or more source(s) 270. The UI 200 may be a conventional personal computer and the information source(s) 270 may be any type of device or component that supplies the desired data or information in a readable form to the UI 200. The archive system may be connected with network components 280, such as workstations, UIs, information sources and other network compatible devices. For example, when the desired data is in the medical environment, the information source(s) 270 may be a readable compact disk (CD) of patient reports or image data (such as CT, MR, CR, X-Ray Angiography and rf, PET, NM, US for example) images from scanners, X-Ray machines and the like.

The UI 200 communicates with a memory system 210 and a media recorder 260 for storing and recording the information from the information source(s) 270. The memory system 210 may be a hard disk, redundant array of independent disks (RAID), or external memory devices that are supported by the UI 200. The memory system 210 provides on-line storage for the initial creation of information file based on the data transfer from the information source(s) 270. The on-line information stored in the memory system 210 may be later accessed by the UI 200 based on universal naming convention (UNC) path of the memory system 210. The media recorder 260 is dependent upon the type of archival storage media that is



used. When DVDs are used as the archival storage media, the media recorder 260 is a DVD recorder that is supported by the UI 200. After the DVDs are recorded, they may be physically placed in jukebox storage 220 as represented by the dashed line.

Due to the storage limitations of the memory system 210, a jukebox storage 220 is provided for retaining selected ones of the archival storage media. The UI 200 supports and communicates with the jukebox storage. Before the memory system 210 reaches its storage capacity, the media recorder 260 records one or more information groups onto at least one archival storage medium, such as a DVD. After the DVDs are recorded, they are mounted in the jukebox storage 220 such that the information on the DVDs are in "near-line" storage. Thereafter, the information on a DVD mounted in the jukebox storage 220 can be accessed by the UI 200 based on the UNC path of the DVD within the jukebox storage 220. One or more jukebox storage units may be incorporated into a single archive system.

Information may be present in the on-line and near-line state at the same time. When near-line information is deleted, the near-line state ends for that particular information. However, such near-line information will not be deleted but will receive an "end" date when the information is no longer valid.

Archive storage 230 is provided for off-line storage of the DVDs. The archive storage 230 may be shelf-type storage remote from the rest of the system. The archive storage 230 may include either duplicate copies of DVDs

mounted in the jukebox storage 220 so that the information appears in both near-line and off-line states or DVDs that are removed from the jukebox storage 220 and are present only in the archive storage 230. The DVDs  
5 are physically placed in the archive storage from either the jukebox storage 220 or the media recorder 260 as represented by the dashed lines. Alternatively, some form of robotic manipulator may be employed to exchange archival storage media between "near-line" and "off-line"  
10 status. When the information contained on an off-line DVD is attempted to be accessed by the UI 200, a volume label of the DVD that was removed or copied from the jukebox storage 220 will result, indicating its "off-line" status.

15 The present archive system allows for a virtual archive 100 to be provided as illustrated in Fig. 1(c). The virtual archive 100 may include a set 110 of information groups 110A, 110B, and 110C that are in combinations of on-line, near-line, and off-line states.  
20 In one example, information group 110A is stored on-line and information groups 110B and 110C are stored near-line. In this case, a primary database 120 is stored in a separate medium from the other information groups 110A, 110B, and 110C of the virtual archive 100. Within each  
25 of the information groups 110A, 110B, and 110C, local databases 112A, 112B, and 112C are respectively stored. The primary database 120 has a database file for the constituent information groups of the virtual archive 100 to all of the local databases 112A, 112B, and 112C, the  
30 status of the storage media, the recorded time of the

storage media, the type of data in the storage media, and other descriptive references relating to this virtual archive 100. Thereby, all of the related information groups 110A, 110B, and 110C can be readily accessed and retrieved despite being retained in different storage states.

The virtual archive 100 may be merged and modified as illustrated in Fig. 1(d). In this example, two information groups 110A and 110B may be retrieved from the virtual archive 100 to create another virtual archive 150. A set 160 of information groups 160A and 160B is created as duplicates or subsets from the information groups 110A and 110B. As above, the virtual archive 150 includes a primary database 170 for relating local databases 162A and 162B contained within the information groups 160A and 160B and retained in different storage states. Accordingly, the virtual archive 150 includes the information groups 160A and 160B while the information groups 110A, 110B, and 110C in the virtual archive 100 remain intact.

When the media recorder 260 records a DVD, a unique identifier is encoded thereon. This identifier uniquely identifies each DVD so that a DVD can be tracked, managed, and interchanged at different locations. Thereby, DVDs from one archive system can be transparently accessed and/or merged with DVDs from another archive system. In one embodiment, the unique identifier is a concatenation of a distinctly assigned volume label followed by values representing the recorded year, day within the year, hour, minute within the hour,

seconds, and milliseconds. For example, a DVD assigned a volume label of 1.2.840.113815 on September 22, 1999 at 9:12:46.157 AM would have the following identifier:

1.2.840.113815.1999265091246157.

5 In this example, no two primary DVD copies will have the same identifier. The identifier is machine-readable and conforms to the universal disk format (UDF) standard volume label format for removable media.

Each DVD also includes a self-contained database  
10 file for holding all of the meta-data required to completely describe a procedure or study stored on that DVD. For instance, in the medical environment, the self-contained database may include all of the demographics for a patient required to adequately review the clinical  
15 procedures, stored on that medium. Digital Image Communications for Medicine (DICOM-3) is preferably used to implement this database file. DICOM-3 is a known standard for enhancing the ability of medical imaging devices and equipment to transfer medical images and  
20 information between systems, such as between a computer tomography (CT) scanner, a workstation, and a printer.

By requiring each information group or the information for each patient to be entirely contained on one storage medium or DVD, it is possible to utilize the  
25 self-contained database file for independently accessing, viewing, and processing each DVD. This addresses a deficiency associated with conventional HSM based systems using DICOM-3, with which patient data often spans two media, such as two magnetic tape units. Each DVD may be  
30 exchanged between archive systems without secondary

database transactions to fully describe a procedure or study. Also, standard clinical imaging stations may directly read the DVDs, allowing review of archived images outside of the archive system in which it was originally created. This is in contrast to prior art systems in which the contents of one storage mechanism such as a magnetic tape are meaningless outside the context of the archive in which it was created. In the event of failure of the primary archive system according to the present invention, information from the DVD may be retrieved on a different archive system, computer, or imaging station.

System auditor data is also recorded on the DVDs for allowing information to be mounted in or removed from an archive system. Information from the DVD's comprising the archive is automatically synchronized to the primary information database whenever information is added to or read from a DVD. The system auditor allows DVDs from different archives to be combined to create a new archive with a single coherent primary database. When a DVD is removed from the system, the system auditor changes the location and status information for that DVD in the primary database to the off-line.

The archive system may also include software for automated network-based media duplication. This duplication software allows for the creation of exact duplicates of the DVDs stored in the archive system on a network attached thereto. Thereby, exact copies of primary archival storage media may be created and stored at different physical locations so that a complete

restoration of the archived information may be performed in the event of a disaster. When using the duplication software in combination with the system auditor, an empty primary database may be completely rebuilt from  
5 information stored on archived storage media. If an archive system is completely destroyed, duplicate off-line media can be used to completely reproduce the original information system on a new archive system.

An embedded image player is also recorded on the  
10 DVDs in this example. The embedded image player includes a computer program for initiating the display of the meta-data, image data, and associated test results contained on the DVD. Thus, the patient data contained on the DVDs may be accessed and reviewed at any later  
15 time without specialized equipment or software.

The archive system selects information to be archived based on predefined system settings. These settings may include: archive oldest or newest first; the number of archival storage media to be written; media  
20 storage capacity; run time; media information staging directory; whether to write to the media or location to move files; and UNC path for study to identifier text file for tracking the media.

In one embodiment of the present archive system, the  
25 archiving process runs as a Windows NT service from the UI 200. The archiving process is configured to run without user interaction with provisions for manually starting, stopping, pausing, and continuing through the UI 200. An application may be provided to update the  
30 status of the archiving process. Fig. 3 is a flow chart

illustrating an example of an archive session according to an embodiment of the present invention. The archive session is initiated at step 300. A search and compilation of studies or procedures are conducted at  
5 step 310. After the studies or procedures are obtained, the number of archival storage media or DVDs needed for these studies or procedures is established at step 320. The studies and/or procedures are staged and segmented into information groups at step 330 and any images are  
10 created at step 340. At step 350, the information groups are written on each DVD.

After writing one or more information groups on a DVD, that DVD is verified at step 360 as to whether the writing has been successfully completed. For example,  
15 verification determines whether each intended information group is contained within the DVD. If the DVD is not verified, it is determined whether writing to the DVD is to be retried at step 364. If a retry is desired or appropriate, writing to the DVD is repeated at step 350.  
20 If the DVD is verified or if a retry is not desired, the information group locations are updated at step 370. Next at step 380, the identifier and index files are recorded before the archive session ends at step 390.

The archive session may be configured to start  
25 archiving at a specified time each day. When the archiving process starts, a number of specified DVDs are written as long as information is available based on a selection process in one embodiment of the present invention. As illustrated in Fig. 4, the present archive  
30 selection is initiated at step 400. At step 410, it is

determined whether any information groups are available. One information group is selected at step 420 when it is determined to be available. At step 430, it is determined whether the selected information group has  
5 already been archived. If the selected information group has not been archived, it is determined whether this information group was created after an archive time setting has expired. This archive time setting is a predefined value corresponding to a desired minimum time  
10 to age before archiving is performed.

If the selected information group was created before the archive time setting expired or was previously archived, the process increments to the next information group, if available, at step 425 without  
15 storing the information group for further processing. If the selected information group was created after the archive time setting expired and was not previously archived, this information group is retained for further processing at step 445.

When no further information groups are determined to be available at step 410, the archiving order is determined at step 450. If the archiving process is configured to archive the oldest information groups first, the information groups are sorted in ascending  
20 order by date and time at step 453. Otherwise, the archiving process is configured to archive the oldest information groups last and the information groups are sorted in descending order by date and time at step 456. Other metrics may be used to define the archive order.



The first of the sorted information group is selected at step 458 and the amount of available storage capacity for a selected DVD is compared to the selected information group. The available storage capacity is  
5 determined by using a predefined percentage full value for the DVD. The amount of storage capacity remaining if the information group were to be recorded onto the DVD is compared to the percentage full value. The  
10 information group is recorded onto the DVD at step 470 if the percentage full value of the DVD will not be exceeded. However, if recording the information group onto the selected DVD would exceed the percentage full value, another DVD is selected at step 464 and then the available storage capacity of the new DVD is determined.  
15 Accordingly, the information group will not span across two or more DVDs and each information group will be entirely contained within one DVD.

Verification of writing to the DVD is performed at step 480. If writing of the information group is not  
20 verified, it is determined whether a retry of writing the information group is desired at step 482. If the information group recorded onto the DVD is verified at step 480, the process attempts to increment to the next information group at step 484 and determines whether  
25 another information group is available at step 490. If another information group is available, the next information group is selected at step 458. However, if no more information groups are determined to be available at step 490, the archive selection ends at  
30 step 499.

100-443886-100